

ESP Kick-Off Workshop Introduction

Plan, legalities, introductions, overview

Presenter: Tim Williams (ALCF)

October 18, 2010



Legalities and Security

- **Workshop open only to registered participants**
 - Your name badge is your pass
- **IBM AECl agreements**
 - Agreement for Exchange of Confidential Information
 - Each individual's home institution must have signed
 - Base agreement
 - Blue Gene/Q Supplement
 - Confidentiality of Blue Gene/Q information
 - Unsure if information is confidential? Email earlyscience@alcf.anl.gov
- **Please treat Workshop materials, discussions as confidential**
 - Next-generation Blue Gene/Q hardware specifics
 - Software for Blue Gene/Q

AECl signed with 30 institutions



Who's Here?

- **77 participants from 23 institutions**
- **Representatives from Early Science Program projects**
 - 39 People, 16 institutions
- **7 IBM representatives**
- **12 ALCF Catalysts & Performance Engineers**
 - Applications/technical questions
- **Libraries, Tools, and Programming Models Project participants**
 - Kalyan Kumaran (Kumar)
 - 26 People, 9 institutions
- **Robert Scott (ALCF User Services)**
 - Accounts, login, CRYPTOCards
- **Sue Gregurich (ALCF)**
 - Workshop logistics



Plan for Workshop

- **IBM talks on Blue Gene/Q**
 - Hardware
 - Software
 - Performance, simulator
 - Compilers
 - Messaging
- **GFDL preliminary performance results**
- **15-minute ESP project presentations**
 - Current code, methods, parallelism, dependencies
 - Needs for Blue Gene/Q
 - Plan for next 6 months
- **From ALCF/MCS**
 - Performance tools plans
 - Programming models
 - Mira acceptance-test applications
 - ALCF I/O plans
 - BG/Q Tools Project
- **Breakout Session (Tools project)**
 - Performance monitors update
- **Working Lunches**
 - Help with ALCF resources
 - Discuss ESP postdocs
 - Plan future workshops
- **RECEPTION TONIGHT**

Group Photo – Monday 3:30 Coffee Break



Early Science Program

- **ALCF Next-Generation IBM Blue Gene/Q machine *Mira***
 - Delivered 2012
 - Production 2013
- ***N Months between acceptance and production:***
Early Science Period
 - 2 billion core-hours dedicated to ESP Projects
 - 12-15 projects → large blocks computer time
 - Must burn time fast
 - 2 years to get ready
 - Mira has ~5X cores of present-day Blue Gene/P machine Intrepid
 - » 20X as fast (10 petaFLOPS)
 - Prepare applications to scale
 - » Node-level parallelism



ESP Projects

- **Call for proposals April 2010**
- **40 viable proposals submitted**
- **16 awards**
 - Target usage on Mira (50 – 150 million core hours)
 - Development time on Intrepid (4 – 15 million core hours over 2 years)
 - About 15 postdoctoral positions focusing on ESP projects
 - Collaboration with ALCF staff and IBM



Science Area Representation

| | Awarded Projects | All 40 Proposals |
|-------------------|------------------|------------------|
| Chemistry | 2 | 4 |
| Nuclear Structure | 2 | 2 |
| Biology | 2 | 4 |
| Materials | 1 | 2 |
| Astro/Cosmology | 2 | 8 |
| Climate | 1 | 4 |
| Fusion | 1 | 2 |
| Geophysics | 1 | 1 |
| Combustion | 2 | 2 |
| Plasma | 0 | 2 |
| Energy | 1 | 1 |
| Computer Science | 0 | 2 |
| CFD/Aero | 1 | 5 |
| Atomic Physics | 0 | 1 |



Institutional Representation

| | Awarded Projects | All 40 Proposals |
|---------------|------------------|------------------|
| National Labs | 7 | 14 |
| Universities | 10 | 24 |
| Industry | 0 | 1 |



National Representation

| | Awarded Projects | All 40 Proposals |
|---------------|------------------|------------------|
| United States | 15 | 37 |
| Foreign | 1 | 3 |



Algorithm/Method Representation

| | Awarded Projects | All 40 Proposals |
|-----------------------|------------------|------------------|
| Structured Grids | 8 | 25 |
| Unstructured Grids | 5 | 19 |
| FFT | 6 | 19 |
| Dense Linear Algebra | 3 | 9 |
| Sparse Linear Algebra | 4 | 12 |
| Particles/N-Body | 5 | 11 |
| Monte Carlo | 2 | 2 |



Develop on Present-Day ALCF Machines


- **ALCF's Blue Gene/P machine, *Intrepid***

- 40K nodes, 160K cores
- 3D torus primary interconnect
- 2 GB per node
- 8 PB online disk

- **ESP allocations for *development***

- Prepare applications for Mira

Intrepid Usage by Program



INCITE: Novel Computational Impact on Theory and Experiment

ALCC: ASCR Leadership Computing Challenge

ASCR: DOE Office of Advanced Scientific Computing Research



END



ESP Proposal Solicitation and Review

- **Forthcoming Call for Proposals publicized**
 - SC09 talks at ANL booth by Messina and Riley
 - Flyer mailed to all ALCF users (about 500) on 18 December 2009
- **January 29, 2010 – Call for proposals issued**
 - ALCF website, INCITE website, ANL main website,
 - Email to over 800 people: INCITE and Discretionary ALCF users, 2009 INCITE proposal submitters, BlueGene Consortium, ad hoc list including SciDAC recipients
- **April 29, 2010 – Call for proposals closed**
- **Proposal Review**
 - Computational Readiness: 2-3 ALCF staff reviewers per proposal
 - Science Impact: 1-2 reviewers with science domain expertise (from ALCF, ANL, and larger scientific communities)
 - Ranking process resulting in recommendations
 1. Numerical scores from 2 review phases
 2. Group discussion by ALCF staff reviewers
 3. ALCF and CELS management discussion
- **PIs notified of results August 6, 2010**



Library and Tool Usage

NOT SHOWN: MPI, MPI-IO, OpenMP

| | Recommended | All 40 Proposals |
|----------------------------------|-------------|------------------|
| NetCDF | 3 | 5 |
| HDF | 5 | 11 |
| VisIt | 2 | 6 |
| ParaView | 2 | 5 |
| IDL | 1 | 1 |
| FFTW, FFTPACK5, P3DFFT, ESSL FFT | 5 | 11 |
| BLAS, ESSL BLAS | 4 | 5 |
| LAPACK | 2 | 6 |
| ScaLAPACK | 1 | 3 |
| MASS, MASSV, ESSL | 1 | 3 |
| parMETIS | 1 | 1 |
| PETSc | 1 | 1 |
| Python | 1 | 4 |
| Tau | 0 | 1 |
| Scalasca | 0 | 1 |
| Charm++ | 2 | 3 |
| Chombo | 1 | 1 |
| GASNet | 1 | 1 |

COLORS: BG/Q PORT
Kumar's tool effort
IBM
Other efforts

